U.S. DEPARTMENT OF COMMERCE
National Technical Information Service

AD-A028 722

Preliminary Investigations Concerning the Training of Tactical Decision Making Behavior

Naval Training Equipment Ctr.

July 1976

TECHNICAL REPORT: NAVTRAEQUIPCEN IH-269

PRELIMINARY INVESTIGATIONS CONCERNING THE TRAINING OF TACTICAL DECISION MAKING BEHAVIOR

Human Factors Laboratory
Naval Training Equipment Center
Orlando, Florida 32813

July 1976

Final Report for Period June 1974 - January 1976

DoD Distribution Statement

Approved for public release; distribution unlimited.

TRAINING EQUIPMENT CENTER

NAVTRAEQUIPCEN IH-269

Preliminary Investigations Concerning the Training of Tactical Decision Making Behavior

ROBERT H. AHLERS, JR.
Human Factors Laboratory

July 1976

Approved:

JAMES S. DUVA
Head, Human Factors Laboratory

J. F. HARVEY
Director
Research and Technology Department

NAVAL TRAINING EQUIPMENT CENTER
ORLANDO, FLORIDA
32813

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER IH-269	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Preliminary Investigations Concerning the Training of Tactical Decision Making Behavior		5. TYPE OF REPORT & PERIOD COVERED Final Report June 1974 - Jan 1976
		6. PERFORMING ORG. REPORT NUMBER NAVTRAEQUIPCEN IH-269
7. AUTHOR(s) R. H. Ahlers, Jr.		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Human Factors Laboratory Naval Training Equipment Center Orlando, Florida 32813		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NAVTRAEQUIPCEN Task No. 3751-4
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Training Equipment Center Orlando, Florida 32813		12. REPORT DATE July 1976
		13. NUMBER OF PAGES ~~32~~ 34
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Tactical Decision Making	Bayes Theorum
Training	Wald Model
Command and Control	Adaptive Training
Computer Models	

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

An accelerating trend for military decision making in command and control situations is to provide the decision maker with statistically processed data. There are obvious benefits to be derived from training a decision maker to be a more efficient user of such diagnostic data. But there is little empirical evidence that training is effective in bringing about an enhancement of decision making performance.

(Cont'd)

PRICES SUBJECT TO CHANGE

Block 20 continued

Two experiments are reported in which it was attempted to determine the effectiveness of a scenario approach for training individuals to make an abstract type of tactical decision based upon probabilistic data. Two questions were of interest in the first experiment: (1) Can appropriate decision making behavior be shaped without providing specific training in the underlying statistical principles? and, (2) Can an adaptive training procedure be successfully utilized in the training of a cognitive skill such as decision making? The second experiment was designed to evaluate a technique of providing performance feedback in order to maintain subject motivation.

It was found that decision making behavior can be shaped without providing explicit training in the underlying statistical principles. An automated, adaptive procedure offered certain advantages for structuring the training session. The performance feedback which was inherent in the adaptive model appeared to provide sufficient motivational cues.

PREFACE

The author would like to express appreciation to John Hearn who conducted the experimental sessions and who contributed significantly to the software development.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

SECTION I

INTRODUCTION

An accelerating trend for military decision making in command and control situations is to provide the decision maker with statistically processed data. This is particularly true where noisy or degraded data, such as sonar bearing information, are to be dealt with. For example, the MK-113 Mod 10 Submarine Fire Control system presents highly processed sonar bearing data to the commanding officer (CO)/approach officer (AO) via the MK 81 Commanding Officer Tactical Display. In order to effectively interpret such data, the user must be aware that they represent only a sample of a particular environmental state and are therefore fallible. The primary task of the decision maker is to reconstruct the environment, i.e., determine the true geographic and temporal relationships between own ship and the target, by optimally evaluating the available relevant data.

The basic problem is therefore one of statistical inference. This general area has been studied for some time, and certain fundamental procedures for optimizing the predictive value of data have been established.

Also well established is the observation that an untrained individual does not optimally evaluate data.(e.g., Snapper & Peterson, 1971). A number of investigators have found that individuals tend to require more data than necessary to reach certain types of decisions. This has been interpreted as support for the notion that information is incompletely extracted from each piece of data, and that, therefore, the subject must acquire a larger data sample than should have been necessary if each datum were used efficiently.

The benefits are obvious for training a decision maker to be a more efficient user of diagnostic data. In an operational setting there is always a cost associated with acquiring additional data. For example, the risk of counter-detection is a cost associated with acquiring more data in a submarine tactical encounter. An encounter should proceed as rapidly as possible from the target acquisition phase to the final attack phase in order to minimize counter-detection. A prerequisite for entering the attack phase is knowledge of the projected target track. Such an estimate of a future event is inherently probabilistic. Since the projected track is derived primarily from the bearing track history, it is apparent that the certainty of the projected track would be directly related to the information that can be extracted from the history. The AO is faced with the tradeoff between waiting for additional bearings, thereby increasing the risk of counter-detection, and attacking without sufficient knowledge of the target's future position. By training the AO to extract a greater amount of information from each bearing, the number of sonar bearings could be reduced without affecting the quality of the decision. The attack phase could be entered more rapidly, thereby minimizing the risk of counter-detection.

While the above argument appears to be valid, there is little empirical evidence that training can induce an enhancement of decision making performance. A vast body of literature exists that concerns the general problem of decision making based upon statistical inference, but there has been

an apparent lack of effort directed specifically at researching the <u>training</u> of decision making. (Nickerson & Feehrer, 1975). Before an operational decision making training system can be specified, further research must be conducted to establish training principles and procedures for effective decision-making skills acquisition.

In general, two approaches to training may be taken. One, the more traditional, would concentrate on teaching the abstract fundamentals of statistical inference. For example, it may be taught that the ability of a data sample to accurately reflect a population parameter is proportional to the $\sqrt{N}$, where N is equal to the sample size. The second approach would allow the trainee to interact with various abstracted situations in which a decision is required. This method aims at shaping the individual's behavior without actually providing an explicit intellectual rationale to support that behavior. The scenario approach might attempt to train the above principle by presenting a series of decision problems in which the trainee must decide between a set of alternatives on the basis of data sampled from them. By properly sequencing the problems and through application of feedback a trainee should acquire an intuitive feel for the relationship between sample size and sample diagnosticity.

One distinct advantage of the scenario approach to training decision making is that the instructional materials can be designed to closely simulate the conditions under which operational decisions will later have to be made. The degree of transfer from this type of learning environment would be expected to be higher than from the situation in which the trainee is exposed to primarily a classroom-oriented curriculum. In addition, relevant aspects of an operational problem may be selectively emphasized in an attempt to eliminate certain behavioral deficiencies which the trainee may be exhibiting. The scenario approach could attack such deficiencies directly by forcing the trainee to repeat appropriate selected problems until he is able to perform at an acceptable level.

The above justifications for using the scenario approach imply that performance measurement techniques must be an integral part of the training environment. This, of course, should also be the case for the more traditional approach. However, the techniques developed for the scenario approach may be utilized for more than evaluating trainee performance and directing the training process -- they may be directly applied to the operational setting for the purpose of assessing the training effectiveness of the entire training system. This information may then be fed back into the system in the form of modifications or enhancements so that training effectiveness might be increased.

A training system such as that outlined above, involves the transmission of vast amounts of information, i.e., scenarios must be generated and displayed dynamically to the trainee, the trainee's performance must be assessed and feedback provided to him, records must be kept, and, ideally, the training curriculum should be structured so as to minimize instructor intervention in those areas which may be more efficiently handled by other means. In order to execute these functions an automated computer-based system would be the most effective vehicle for scenario presentation and training management.

It can be recognized that there are strong arguments to support a training system based upon a scenario approach for training decision making behavior in a tactical context. It should be appreciated, however, that a relatively small set of research studies may be used to support the development of such a system. In order to provide the necessary research support for such a training system, a series of in-house research studies has been initiated.

A number of basic issues need to be investigated before a prototype decision making training system can be designed. Several of those questions which could be investigated in an abstract decision making situation were selected for the present research. Two initial experiments have been conducted and are reported on separately below.

SECTION II

EXPERIMENT I

STATEMENT OF THE PROBLEM

The first experiment directly addressed the problem of determining the effectiveness of a scenario approach for training individuals to make an abstract type of tactical decision based upon probabilistic data. Two questions were of interest: (1) Can appropriate decision making behavior be shaped without providing specific training in the underlying statistical principles? and, (2) Can an adaptive training procedure be successfully utilized in the training of a cognitive skill such as decision making? It was felt that these issues would be of particular interest if the functions of instructional material sequencing and performance evaluation and feedback were to be automated, as in a computer-driven training system.

METHOD

EQUIPMENT. The experiment was conducted using the Human Factors Laboratory's Automated Display Controller System (ADCONS). This facility consists of a PDP-9 with 32K words of 18 bit memory, two Cathode Ray Tube (CRT) refresh-type displays with lightpens, disk and DECTAPE mass storage devices, and a control teletype. The entire system is more complex, but only the hardware described above was utilized for the present experiment.

The computer program that was written for the experiment served several major functions. Among them: (1) It contained scenario generation logic to present a graphic display of each problem to the trainee and to accept his responses as indicated by his manipulation of the lightpen. (2) It automatically evaluated the trainee's performance in real-time against an "optimal" model which was resident in the program. Whenever additional data was selected, the model evaluated the accumulated data sample to determine if its acceptance or rejection criterion was exceeded. In this manner, the model controlled the sequence of events within each problem. (3) Three higher level models had the capability of structuring the problem sequence. One contained a linear adaptive logic which changed problem difficulty in response to the trainee's performance. The other two scheduled the problems according to predetermined sequences, described below.

Problems were presented to the trainee on a remote CRT in an experimental room and were simultaneously displayed on a CRT at the control console so that the trainee's performance could be monitored by the experimenter.

SUBJECTS. Male subjects were recruited from the undergraduate curricula at Florida Technological University. Eight subjects were assigned to each of the three training conditions.

PROCEDURE. The experiment was designed to investigate two basic questions concerning the training of decision making:

a. Can performance be improved in a statistical inference task using standard feedback techniques: and,

b. Can an adaptive training procedure be successfully utilized in the training of a cognitive skill such as decision making?

Three training procedures were investigated, each of which is detailed separately below. Each was designed to structure the presentation of a set of problems involving statistical inference so that learning would be facilitated.

The problems were variations of one basic scenario. Appendix A contains the instructions for the subjects which describe the scenario in detail. The subject was told that each problem required him "to make a decision analogous to that of a submarine officer investigating a report concerning the presence of an enemy submarine." A message presented on the CRT informed the subject of the probability that an enemy submarine was patrolling in his area. See Figure 1. The subject's task was to evaluate this "intelligence report" based upon supplementary data and then indicate the presence or absence of the enemy submarine when he was "fairly certain" of his decision. The supplementary data were sampled one point at a time and indicated either the presence or absence of the enemy submarine depending on sampling bias.

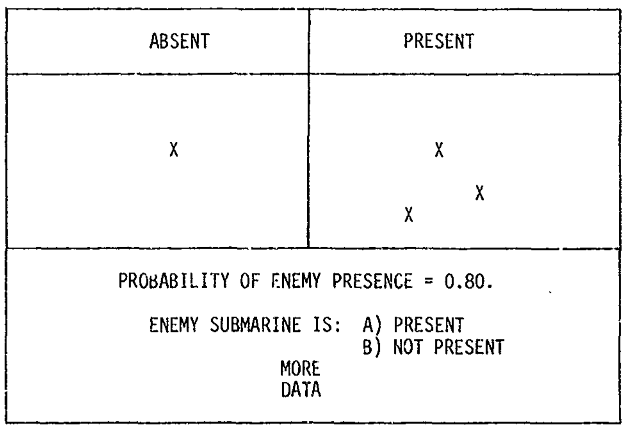

Figure 1. Scenario CRT display as presented to the trainee

A modified version of the optional stopping model of Wald (1947) was used to monitor the data acquisition process of the subject in real-time. The model was originally developed for industrial inspection applications and defines a procedure for the efficient estimation of output quality based upon the information contained in small samples. This procedure is

functionally similar to that followed in an intelligence gathering situation, i.e., a new piece of information is collected, the situation is reevaluated and, if no decision can be made more information is collected.

Analysis of preliminary work in this laboratory using the Wald model had indicated that subjects do not learn to observe the rigorous sampling criteria of the model. For this reason the model was modified to include tolerance intervals around each of the criteria, such that a correct decision could be reached whenever sample composition was within one unit of a cutoff. Allowing the model to tolerate a small degree of error did not affect the salient statistical relationships that were being trained, but it had the effect of increasing the number of correct responses made by the subject, thereby increasing his motivation. It was felt that maintaining an "optimum" model was of secondary importance to having a slightly degraded one which was better suited for training.

The model was set up to allow the subject to reach a terminal decision concerning the presence of the enemy when there was a 90 percent chance of his making the correct response. This 90 percent confidence interval remained constant for all problems. Following the terminal decision of each problem, a subject was presented with an informative feedback message on the CRT. One of four messages appeared:

a. Your response was correct!

b. An enemy submarine was not present.

c. An enemy submarine was present.

d. You did not have sufficient data.

Problem difficulty was manipulated within the training session such that the problems tended to become more difficult as the session progressed. Problem difficulty was defined in terms of the a priori probabilities, i.e., a high a priori probability of enemy presence resulted in a problem which required little supplementary data to solve and was therefore considered to be less difficult than one with a lower initial probability of enemy presence.

The sampling bias for the supplementary data reflected the probabilities associated with the initial intelligence report, i.e., the supplementary data tended to confirm the original report, and the strength of this tendency was a direct function of the magnitude of the a priori probability of enemy presence. For example, if the presence of an enemy submarine was reported as being 90 percent probable, 90 percent of the supplementary information would confirm the report.

As stated previously, the problem difficulties were varied throughout a training session according to one of three schedules, corresponding to the three experimental groups:

a. Adaptive group. Problem difficulty was varied as a function of past performance, with the a priori probability of enemy presence functioning

as the adaptive variable. For each difficulty level the progressive criteria was set at four out of the last five trials correct; the regressive criteria was slightly less stringent causing the problem difficulty to regress when a subject scored three out of five incorrect trials. The enemy probability for the initial problems in the training session was .74 and changed in steps of .02. Each change in problem difficulty was signaled to the subject with an appropriate message during the intertrial interval. The two alternate messages were:

(1) YOUR PERFORMANCE HAS DETERIORATED. CONSEQUENTLY, YOU WILL BE GIVEN A GROUP OF REMEDIAL TRIALS.

(2) CONGRATULATIONS! YOU HAVE BEEN PERFORMING VERY WELL. THE FOLLOWING PROBLEMS WILL BE MORE CHALLENGING TO BETTER MATCH YOUR ABILITY.

b. Self Adaptive group. Each subject was allowed to choose the difficulty of the next problem during the intertrial interval by responding to one of the following alternatives with the lightpen:

I WOULD LIKE THE FOLLOWING TRIALS TO BE: A. MORE DIFFICULT
B. LESS DIFFICULT

Problem difficulty was shifted in the appropriate direction by changing the prior probability of enemy presence by .02. If the subject chose neither alternative, the problem difficulty would remain constant.

c. Fixed Progression group. Problem difficulty for the subjects in this group was not determined by the individual subject, as in the other two groups, but was based on the Adaptive group's performance. The modal difficulty for each problem in sequence was used as the difficulty for the same sequential problem in the Fixed Progression group. So, although the problems did not adapt individually, they could be considered to vary according to a "group adaptive" scheme in which the combined problem-by-problem performance of one group of subjects determined the sequence for the Fixed-Progression group.

For all groups, the entire session was divided into three phases, Pre-training test, Training, and Post-training test. Eight problems were given during Pre-training, all with a prior probability of .62. This represented the criterion difficulty level, i.e., the problem difficutly with which the subject was to be trained to deal. Following this set of Pre-training problems, the supplementary instructions for the appropriate group were read to the subject. The training session then began, and training continued until the subject progressed to the criterion level problems. This represented the start of the Post-training test phase which lasted until the subject completed six problems at the criterion level of difficulty.

RESULTS

Data were analyzed from the last six Pre-training problems, the first two being considered as familiarization trials. These data were contrasted with data obtained from the six Post-training problems for an analysis of training effects.

Two dependent measures were of interest - the number of correct terminal decisions and the amount of data sampled before reaching a terminal decision in each of the correct problems. A decision was considered to be correct if the subject's choice of underlying distribution matched that of the model. It is important to point out that since the sampled data were fallible, the model's choice of distribution was not always the correct one. Therefore, the dependent measures reflected how well the subject's performance matched that of an optimal decision maker, i.e., the model, and not how correct his decisions were, based upon the real world situation. In other words, his terminal decisions per se were not of interest, but his decision process, the manner in which he arrived at those decisions, was the primary concern. The decision process is the aspect reflected in the performance measures which is amenable to training and so was emphasized.

Each of the measures, number of correct decisions and sample size, was analyzed in a 3 x ? split plot design with the three types of training procedures as the between subjects treatment and Pre- or Post-training as the within subjects variable. Analyses were run on a PDP-9 using a split plot ANOVA program written in FOCAL (Breaux, 1972). Tables 1 and 2 contain the summary ANOVA data for each of the measures.

TABLE 1. ANOVA SUMMARY TABLE FOR THE NUMBER OF CORRECT TERMINAL DECISIONS

Source of Variance	dF	MS	F
Training Procedure (A)	2	7.15	2.65
Pre/Post-Training (B)	1	18.75	13.07 *
A x B	2	3.06	2.13
Subjects Within A	21	2.70	
Subjects Within A x B	21	1.43	
Total	47		

* $p < .002$

Examining the effects on the number of correct terminal decisions (see Figure 2), it can be seen that the significant main effect, Pre/Post-training, shown in Table 1, was found under all three training techniques. The performance improvement was apparently greatest in those subjects who had undergone the adaptive training procedure -- an increment of over 100 per cent.

TABLE 2. ANOVA SUMMARY TABLE FOR AMOUNT OF DATA SAMPLED ON CORRECT TRIALS

Source of Variance	dF	MS	F
Training Procedure (A)	2	187.17	4.72 **
Pre/Post-Training (B)	1	168.04	4.45 *
A x B	2	144.99	3.84 *
Subjects Within A	21	39.63	
Subjects Within A x B	21	37.76	
Total	47		

* $p < .05$
** $p < .02$

The first order interaction of the training procedure with Pre/Post-training only approached, but did not achieve, significance at the .05 level. Therefore, no rigorous statement may be made concerning the efficacy of a particular training technique. However, the analysis does support the overall hypothesis that individuals may be trained to become better decision makers without providing specific training in the underlying statistical principles. Further support is given by the finding that of the 24 subjects trained in the experiment, only three performed more poorly on the Post-training test than on the Pre-test; none of these was in the Adaptive Group.

The other measure of decision making skill, amount of information sampled, is graphed in Figure 3. It had been expected that this measure would decrease with training as the subjects learned to overcome their expected conservatism. The opposite effect was found. The average sample size before training was 12.75 data points; after training, 16.5. This effect was significant at the .05 level. This finding is misleading, however, unless considered in the light of the first order Training Procedure X Pre-/Post-training interaction, also significant at the .05 level. Figure 4 graphically represents this interaction. The largest difference between Pre- and Post-training sample sizes was found for the Self-Adaptive group. Using the Scheffe' test (Hays, 1963) with a 95 percent confidence interval, the only significant Pre-/Post-training contrast was for the Self-Adaptive group. Training using the two alternate procedures did not affect the sampling behavior. A Scheffe' test showed that the significant training procedure main effect was also a result of the conservative sampling behavior in the Post-training phase of the Self-Adaptive group. The only significant contrast was between the Self-Adaptive and Fixed Progression groups.

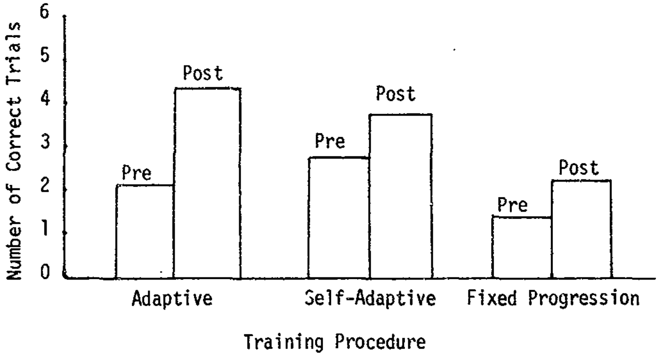

Figure 2. Number of correct terminal decisions during Pre- and Post-training phases as a result of type of training

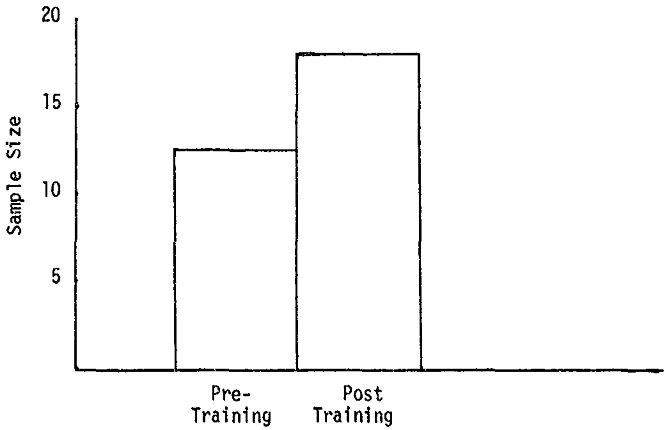

Figure 3. Amount of data sampled during Pre- and Post-training phases

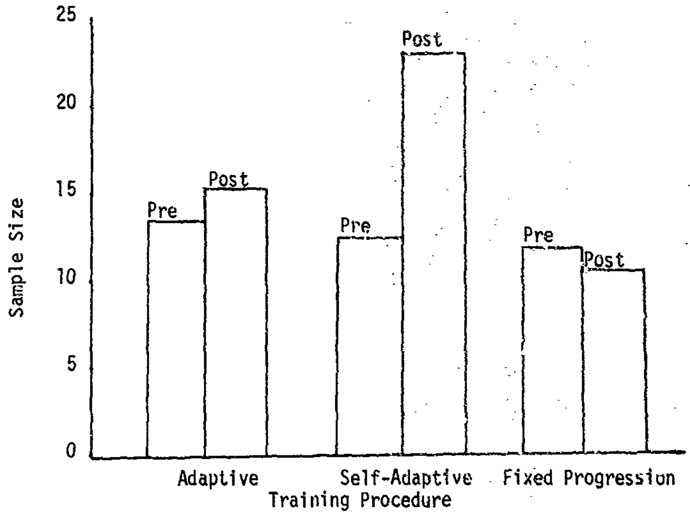

Figure 4. Amount of data sampled during Pre- and Post-training phases as a result of type of training

For subjects in all groups, there was a tendency to select too few data points during the Pre-training phase. Evidence for this behavior can be found by looking at the types of errors committed during this phase. Only 17 percent of the errors involved requesting more data than was necessary to reach a terminal decision while 82 percent were the result of making a decision based upon insufficient data. (The other one percent was due to choosing the incorrect alternative.) Subjects were, therefore, selecting close to the minimum sample size from the very beginning; training could not be expected to bring about a further reduction.

A point of interest is the change in the sampling behavior of the Self-Adaptive subjects. Unfortunately, there is nothing in the data which would explain the observed conservatism in their Post-training trials. Based upon this observation, it is sufficient to conclude for the purpose of the present experiment that the Self-Adaptive technique is inferior to the other two, given that selecting more data than needed is an undesirable behavioral characteristic.

DISCUSSION

The data support the hypothesis that decision making behavior can be shaped without providing explicit training in the underlying statistical principles. Further, it was found that an automated adaptive procedure may successfully be utilized to sequence the training scenarios.

It should be recognized that the particular adaptive logic used in the present experiment is most likely not an optimum one. It would be expected that a more appropriate logic would yield an increased training benefit.

It was not the purpose of the present research to attempt to determine such an optimum logic. The determination is largely an empirical problem, and should, therefore, be investigated under conditions closely approximating those under which the training logic would later be utilized. Both the context under which training is to take place and the entrance characteristics of the trainee population would be important considerations. Although the problem was not addressed by the present research for the above reasons, considerable effort should be expended to uncover an efficient adaptive training logic before applying an adaptive automated technique to an operational decision making training context.

There is some question about whether the Wald model accurately described the decision situation which was presented by the scenario information. For example, the Wald model does not consider the impact of the prior probabilities associated with the presence or absence of an enemy submarine, only the sample composition is evaluated to reach a terminal decision. The subjects, on the other hand, appeared to make use of the a priori data, the intelligence report, during the Pre-training phase. (Their behavior during the Pre-training phase is used as a point of comparison since they were behaving naively with respect to the model during this period of performance.) Looking at the data collected during this phase, it can be seen that the subjects were frequently able to correctly "outguess" the model. This would suggest that the model, or models, which the subjects were following was more optimum than the prescriptive model used as the basis for evaluating their performance. In fact, on those trials in which the subjects reached a terminal decision before the model would have allowed, the decisions were correct 76.4 percent of the time. Further, on those trials, all subjects except one made more correct decisions than incorrect ones.

These observations point to the conclusion that the Wald model inappropriately described the decision situation which was presented by the scenarios. Because the model did not consider the prior probabilities, it forced the subjects to aggregate more information than should have been required to operate within the prescribed error tolerances. The model was thus forcing the subjects to behave in a more conservative manner than was necessary.

This interpretation of the data does not invalidate the implications for the training of decision making found in this experiment. In fact, it may be argued that since decision making behavior could be shaped to approximate a nonappropriate model, the use of a more appropriate model should yield enhanced training effects and higher levels of decision-making performance.

SECTION III

EXPERIMENT II

STATEMENT OF THE PROBLEM

The second experiment was designed to evaluate a technique of providing performance feedback in order to maintain subject motivation. Based upon opinions elicited from a high percentage of Experiment I subjects during informal exit interviews, it was observed that many subjects had "lost interest" in the task at some point during the session. This was not unexpected, due to the high degree of concentration involved and the riskless nature of the decisions involved. These task characteristics are not specific to the task used, but may be expected to be present in any non-real world decision making environment. The typical solution for this problem has been to employ a monetary payoff structure to make the task more interesting from the subject's point of view. While this approach has been satisfactory for laboratory paradigms, its use in an applied training setting would be impractical. It would be desirable to be able to exploit certain features of the training task to serve a similar function.

The present experiment investigated such a potential source of subject motivation. In particular, each subject's prior performance level was updated in real-time and continuously displayed to him in one of two formats. It was felt that each subject, by knowing how well he had been performing, would feel that he was in competition with himself and would, therefore, be motivated to perform to the best of his current level of skill.

METHOD

EQUIPMENT. The hardware utilized in the present experiment was the same as that used for Experiment I. The software differences were significant, but the same general functions were performed.

SUBJECTS. Thirty-six male subjects were recruited from the undergraduate curricula at Florida Technological University. Subjects were randomly assigned to each of the six treatment cells of the experimental design.

PROCEDURE. The same basic scenario from Experiment I was used in the present experiment but slightly different information was presented to the subject, and the scenario was driven by a different model. The instructions to the subjects provide a clear impression of the scenario situation as presented to the subjects (see Appendix B).

A representation of the CRT display is shown in Figure 5. The major formatting change from the display used in Experiment I was an increase in the spatial separation of the response areas. This was done in an attempt to minimize unintentional responses caused by the inadvertent aiming of the lightpen at a response area. This condition occurred infrequently in the previous experiment and, therefore, represented a small source of uncontrolled variance in the data. However, the condition was easy to alleviate and so the appropriate changes were made.

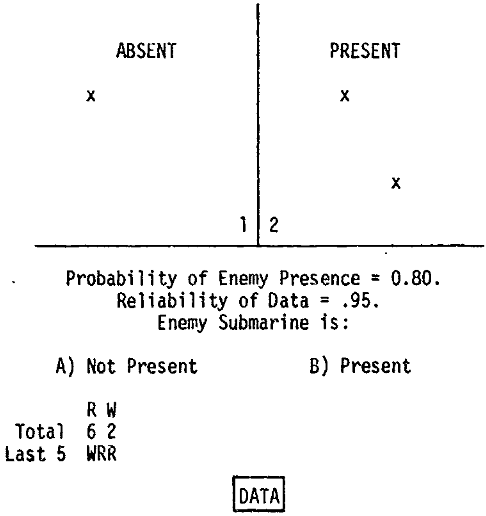

Figure 5. Representative scenario display. Selected features of the display were not presented under certain treatment conditions (see text).

As noted in the prior discussion section, the Wald model was felt to not precisely reflect all the information that was present in the scenario. For this reason a Baysian model (e.g., Hayes, 1963) was selected as the prescriptive model in the present experiment. The model evaluates prior probabilities in the light of the conditional probabilities associated with observed data in order to estimate revised or posterior probabilities.

Bayes theorum may be stated as follows:

$$p(H|d) = \frac{p(H) \times p(d|H)}{p(d)}$$

where d represents the last data point which was sampled, an indication of either Present or Absent, and H stands for the hypothesis which is being tested, "The enemy submarine is present." The slash bar should be read as "given." The quantity p(H) represents the prior probability of enemy presence, or the estimate of enemy presence before the currently selected data point is evaluated. This prior probability is modified by the expression $\frac{p(d|H)}{p(d)}$, which represents the impact of the currently displayed data point. For example, if a "Present" point appears it is more likely that the hypothesis "the enemy submarine is present" is true, i.e., the

probability of a "Present" point occurring given that the "eneny present" hypothesis is true ($p(d|H)$) is higher than if the "enemy absent" hypothesis were true. This conditional probability, $p(d|H)$, is normalized by $p(d)$ which represents the overall probability that a particular class of data point will be ovserved.

The evaluation of the three right-hand terms of the equation will yield a revised probability estimate, the probability that the hypothesis of interest is true given the additional information in the current data point. This revised probability estimate becomes the current estimate of $p(H)$ for the purpose of evaluating the subsequent data point.

Data sampling and the consequent revision of the probability estimate of the hypothesis continues until the estimate exceeds a previously established criterion. In the present experiment, this probability level was set at .99. That is, the subjects were able to correctly decide between the two hypotheses, "enemy absent" and "enemy present" when $p(H|d) = .99$. However, even if the sampling criteria of the model were met, the selection of the indicated hypothesis would be in error 1% of the time. Since the subjects' responses were scored in relation to the model and not to the true state of the world, this source of error was invisible to the subjects.

It is to be expected that individuals do not explicitly evaluate information according to the rigorous procedure specified by Bayes theorum. However, the model is useful because it takes into consideration all information sources, and it does seem to provide a reasonable description of how individuals evaluate information. Further, it is an optimum model in that it presents the most efficient manner in which to evaluate information presented in situations which conform to the constraints of the model.

Three conditions must be met to allow the subject to aggregate information in the manner prescribed by Bayes theorum. First, he must be allowed access to the data sample.

Secondly, he must know the prior probability of one of the two mutually exclusive hypotheses. In this case he was given the initial probability of enemy presence and was to assume that the source of this information was an external intelligence report. In all of the scenarios this probability was set at .80.

Finally, the subject must be given the diagnosticity of the data which he will be requesting ($p(d|H)$). This was presented to the subject as the "reliability of data." It varied from 1.00, high reliability, to a theoretical lower bound of .50, low reliability and of no diagnostic value. That is, for example, a "Present" data point which has a reliability of .50 could occur with equal liklihood if either the "enemy present" or "enemy absent" hypothesis were correct. It would, therefore, be impossible to determine which was actually being sampled from no matter how large a data sample was acquired. On the other hand, a reliability of .90 would specify that 90% of the data points would fall in the category which corresponded to the true hypothesis; discrimination between the two hypotheses would be relatively easy.

Problem difficulty was adjusted in such a manner in the present experiment. The same linear adaptive logic from Experiment I was used to change the reliability of the data in increments of .05. Initially, all subjects started the session solving problems with a data reliability of 1.00. At this level, a correct terminal decision could be made after evaluating only one data point. Scoring four out of five correct decisions at this, or any other level, resulted in the presentation of a message informing the subject that he had been performing very well and that the next series of problems would be more challenging so as to better match his ability. This next series would have a data reliability which was .05 lower.

The regressive criterion was three incorrect out of the preceding five problems. When this criterion was met, a message was displayed which informed him that his performance had deteriorated, and he would be given a group of remedial trials. Data reliability was incremented by .05 for the next series of problems. The session was terminated after 45 minutes.

Three main treatment conditions consisting of different types of motivational feedback were examined. The motivational feedback provided no problem specific information. It was structured to inform the subject how well he had been performing on the problems already completed.

The descriptive feedback presented after each problem was given under all treatments. This post-problem feedback consisted of information which allowed each subject to compare his behavior with that of the prescriptive model. A description of the feedback messages appears in the Procedure section of Experiment I. This feedback was the only source of information which directly impacted on the subject's learning of the task.

One group of subjects served as a control and received no motivational feedback. The two experimental groups differed in the type of motivational feedback given them. The feedback was continuously visible in the lower left-hand quadrant of the CRT and was updated immediately following each terminal decision. The feedback for one group consisted of the running totals of correct and incorrect problems. For the second experimental group, only information on the scoring of the last five trials was provided. Since performance on this set of trials determined the immediate behavior of the adaptive model, it was felt that knowledge of this performance would motivate a subject to "try harder" on those problems which were critical to his advancement in the adaptive sequence. The hypothesized increase in motivation should yield a higher percentage of correct responses and, therefore, a higher difficulty level of the problems at the exit point of the session. A similar but attenuated effect was expected from the group which were given their total summary performance as feedback.

The three feedback treatments were combined orthogonally with two slightly different display formats for the supplementary data. One format was identical to that used in Experiment I; the other contained the same information but, in addition, had a digital readout of the total number of data points in each of the two cells. This was done to suggest a counting strategy and to make that strategy easier to employ. The Bayesian posterior probabilities for any given data reliability and initial prior probabilities

are determined solely by the relative difference between the number of data points falling within the two mutually exclusive classifications. Sample size has no effect on this relationship. It was felt that by suggesting the most efficient strategy to the subject, he would be less likely to search for new strategies to try, and his performance would consequently have a higher internal consistency. Treatment effects between feedback conditions would, therefore, be more obvious.

RESULTS

The performance measures selected for analysis were based upon the number of correct responses. The highest difficulty level attained during the session, the percentage of correct problems, and the correlation between difficulty level and problem number were each analysed in a 3 x 2 factorial ANOVA. (Winer, 1962).

Figures 6 and 7 contain graphs of the average difficulty level by groups as a function of practice. All of the curves show that subjects were able to solve increasingly difficult problems as practice was acquired. The overall correlation between data reliability and problem number was -.88. A T-Test showed this correlation to be highly significant beyond the .001 level, indicating that the level of skill increased as a function of practice. However, the homogeneity of the curves indicates that the various treatment conditions had little effect upon learning. The various F-Tests performed on the data confirm this observation. No significant effects were found in the analyses of the scores reflecting the highest difficulty attained, the percentage of correct problems, or the correlation between difficulty level and problem number.

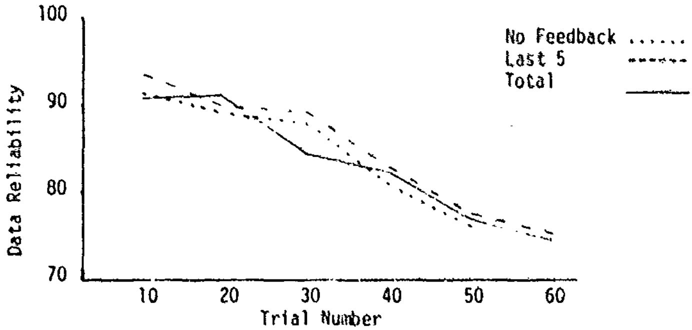

Figure 6. Performance as a function of practice for subjects with digital data readouts

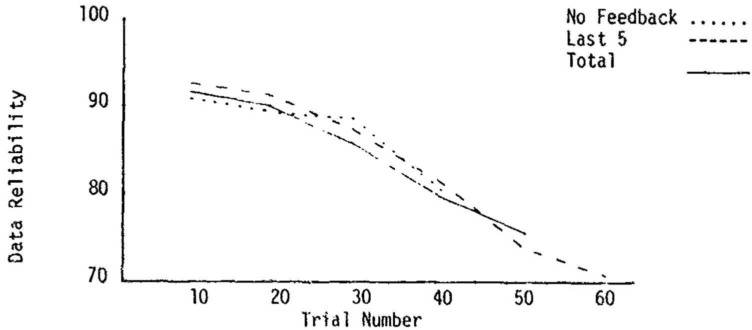

Figure 7. Performance as a function of practice for subjects with normal information display

DISCUSSION

The data analyses indicated that the technique of providing performance feedback was not effective in motivating the subjects to perform better in the type of decision-making training paradigm used. There are several explanations of this finding which may be offered.

Perhaps the subjects who received the feedback did not attend to it. This hypothesis is supported by comments elicited from the subjects concerning their usage of the feedback. Slightly over one-half reported that the feedback was not used or was used only occasionally, generally after several incorrect terminal decisions had been made. A stimulus which was not perceived, i.e , the feedback, in this case, could not be expected to have been instrumental in modifying the subjects' behavior.

The data indicate, however, that the above explanation is not entirely correct. It would be expected that the performance of those subjects who kept track of their performance using the feedback would differ from that of those who did not. This was not found to be true, those subjects who attended to the feedback on every problem did not seem to derive any motivational effects from it. This is an interesting finding, indicating that although these subjects were concerned about their performance, they apparently did not expend additional effort to improve it.

Supporting evidence may be found in a post-hoc analysis of the time spent in extracting information from each data point for those subjects who attended to the feedback. Problems were divided into two categories, critical and noncritical. The former were defined as those problems whose outcome would determine if the next group of problems would be changed in difficulty. A problem which did not meet this definition was considered to be noncritical. If it is assumed that the function of the feedback was to call attention to the critical problems so that subjects would expend

more effort on them, then it would be expected that more time would have been spent evaluating the information before a terminal decision was made. The average latency per data point was 4.372 seconds for the noncritical problems and 3.348 seconds for the critical ones, thereby not supporting the above expectation. The difference was found to be not significant using a matched pairs T-Test. (Hays, 1963). So, even analysis of this relatively fine-grained measure of performance, and including only those subjects who used the feedback, failed to show any differential effects of providing feedback which was designed to be motivational in nature.

An alternate explanation of the lack of significant treatment effects may be found by considering the format of the problem presentation. A salient feature of the adaptive problem sequencing was the presentation of an appropriate message informing the subject of an upcoming change in the problem difficulty level whenever his performance warranted such a change. The messages, along with the display of the current difficulty level, may have been sufficient to allow each subject to keep track of his performance and, thereby, provide sufficient motivation to perform well on each problem. If this were the case, then the distinction between critical and noncritical problems would have been an artificial one, i.e., the subjects might have perceived each problem as contributing equally to the overall performance score.

The average highest difficulty level achieved by the subjects was represented by an information reliability of .66. That is, subjects were able to arrive at correct decisions regarding enemy presence based upon information which was in error 34 percent of the time. Based upon preliminary experience with the problems, it may be stated that this represented extremely good performance. Subjects apparently were strongly motivated to perform well, and the most probable source of this motivation seems to have been the attendant feedback sources present in the adaptive format of problem sequencing.

SECTION IV

SUMMARY

Decision making behavior could be shaped without providing explicit training in the underlying statistical principles. An automated adaptive procedure offered certain advantages for structuring the training session. The performance feedback which was inherent in the adaptive model appeared to supply strong motivational cues. This particular conclusion needs to be investigated further, but it is apparent that the supplementary feedback provided in Experiment II was not an important source of motivation, indicating that the source of motivation was resident in the adaptive structure.

Many of the important questions which need to be answered regarding the application of training principles to a decision making training behavioral objective can only be investigated in a context-specific environment. It is, therefore, recommended that the above fruitful lines of investigation be continued in a more applied context with specific training objectives. Such a setting would lend itself to the investigation of central issues such as the determination of an optimal adaptive logic, the isolation of diagnostic performance measures, and the necessity of providing feedback.

This laboratory is currently utilizing the context of a submarine approach officer command/control task to investigate such questions. A dynamic simulation of an ASW encounter driven by computer models of sonar parameters, ship dynamics, fire-control solution, weapon characteristics, etc. is being used to drive a comprehensive information display. The AO trainee is, thereby, provided with a range of information resources which vary in both relevancy and quality. His task is to draw inferences from these data concerning various target ship parameters. A scenario approach to training is being used, and various alternative feedback techniques and displays are being evaluated. It is anticipated that the results of this current line of investigation, utilizing an applied context, will confirm and extend the conclusions drawn from the experiments reported on above.

REFERENCES

Breaux, R. Analysis of Variance of One, Two, and Three-Treatment Designs for a PDP-8. Behavioral Research Methodology and Instrumentation, 1972, 4(5), 271-272.

Hays, W. L. Statistics. New York, Holt, Rinehart and Winston, 1963.

Nickerson, R. S., & Feehrer, C. E. Decision Making and Training: A Review of Theoretical and Empirical Studies of Decision Making and Their Implications for the Training of Decision Makers. NAVTRAEQUIPCEN 73-C-0128-1, 1975.

Snapper, K. J., & Peterson, C. R. Information Seeking and Data Diagnosticity. Journal of Experimental Psychology, 1971, 87, 429-433.

Wald, A. Sequential Analysis. New York, Wiley, 1947.

Winer, B. J. Statistical Principles in Experimental Design. New York, McGraw-Hill, 1962.

APPENDIX A

SUBMARINE DETECTION SCENARIO BASED UPON WALD MODEL-INSTRUCTIONS FOR SUBJECTS

This experiment will investigate how well you can make a certain type of military decision and how much your performance will improve with practice.

The display you see in front of you is generated by a computer. You will be presented with a series of problems similar in appearance to the one being displayed now. Each problem will require you to make a decision analogous to that of a submarine officer investigating a report concerning the presence of an enemy submarine.

Consider that you have received an intelligence report stating that the possibility exists of an enemy submarine patrolling in your area. (Point out report to subject.) This information is displayed to you in the form of a probability statement, i.e., the probability of enemy presence is equal to .80. This means that there is an 80% chance that a submarine is present.

It is your task to decide if an enemy submarine is in fact present. In order to make this decision, you will be allowed to use your available resources to acquire information relating to the presence or absence of the enemy submarine. Specifically, you will be allowed to interrogate the display to ask for additional data points which will indicate either the presence or absence of an enemy submarine. (Demonstrate). You should be aware of the fact that it is impossible to tell with absolute certainty whether or not an enemy submarine is present. Any one data point, or group of data points, may be erroneous, but a large number of data points will tend to reflect the actual situation. In general, the more data on which you base your decision, the more likely it is to be correct.

It may be useful for you to think of each data point as representing an opinion of an experienced crew member. That is, each time you request more data you are in effect asking if he thinks his equipment is sensing the presence of an enemy submarine. He doesn't know for sure if one is present, but he tells you what he thinks at that particular time. If he indicates that one is present, a data point will appear in the box labeled "Present." If he does not detect a submarine, a data point will appear in the "Absent" box.

You may continue to ask for more data until you feel fairly certain that you can make a correct decision. When this point is reached, use the light pen to indicate your choice. (Demonstrate).

The computer will select the problems presented to you, and it will monitor your performance. If you ask for too much data or make a wrong decision, the computer will display a message indicating the decision that you should have made. If you respond correctly, a message to that effect will be displayed. You will also be informed if you make a decision based upon insufficient data.

Keep in mind that you should choose just enough data to allow you to be fairly certain of your decision. Choosing either too much or too little data will result in the computer scoring the trial as an error.

When you make a selection, or if the computer terminates the trial, the display will go blank and a message will appear informing you of your performance on that trial. This message will then be replaced with the display as it appeared when a choice was made, by you or the computer. You should find both the message and the display helpful for improving your performance on subsequent trials. Following an intertrial interval of five seconds, another problem will be displayed. You will now be given a set of problems. Please try to correctly evaluate the intelligence reports using the minimum number of additional data. It is important that you understand what you are to do so please ask about any section of these instructions which you do not fully understand.

I will be in the next room monitoring your performance and will intercede if you appear to be having any problems. Following this first set of problems I will give you additional instructions concerning the remainder of the session. The first experimental trial will begin shortly.

SUPPLEMENTARY INSTRUCTIONS FOR ADAPTIVE GROUP

The beginning problems of the next set will have a low difficulty level, and you will be able to evaluate the intelligence reports without choosing additional data. The difficulty of the remaining problems will be selected to match your ability. This will be done because we do not want you to waste your time and effort on problems which are far below your capabilities or to struggle with problems which may be too difficult for you. As your performance changes with practice, the problem difficulty will be adjusted to assure that the problems remain challenging without being excessively difficult. You will be informed by a displayed message each time the problem difficulty is to be changed.

Do you have any questions? Again, I will be in the next room monitoring your performance.

SUPPLEMENTARY INSTRUCTIONS FOR SELF-ADAPTIVE GROUP

The beginning problems of the next set will have a low difficulty level, and you will be able to evaluate the intelligence reports without choosing additional data. You will be allowed to change the difficulty of each problem before it is presented. This will be done because we do not want you to waste your time and effort on problems which are too far below your capabilities, or to struggle with problems which may be too difficult for you. For the five seconds immediately preceding each problem the following message will be displayed:

I WOULD LIKE THE FOLLOWING TRIALS TO BE: A. MORE DIFFICULT
B. LESS DIFFICULT

To change the level of difficulty, aim the lightpen at the appropriate alternative and depress the shutter. The unpicked alternative will disappear.

If you do not make a choice, the difficulty level will remain the same. During the course of the experiment, you should try to maintain the highest difficulty level that is possible.

Do you have any questions? Again, I will be in the next room monitoring your performance.

SUPPLEMENTARY INSTRUCTIONS FOR FIXED PROGRESSION GROUP

The beginning problems of the next set will have a low difficulty level, and you will be able to evaluate the intelligence reports without choosing additional data. The difficulty levels of the remaining problems will tend to increase as the session proceeds.

Do you have any questions? Again, I will be in the next room monitoring your performance.

APPENDIX B

SULMARINE DETECTION SCENARIO BASED UPON BAYESIAN MODEL - INSTRUCTIONS FOR SUBJECTS

This experiment will investigate how well you can make a certain type of military decision and how much your performance will improve with practice.

The display you see in front of you is generated by a computer. You will be presented a series of problems similar in appearance to the one being displayed now. Each problem will require you to make a decision analogous to that of a submarine officer investigating a report concerning the presence of an enemy submarine.

Consider that you have received an intelligence report stating that the possibility exists for an enemy submarine to be patrolling in your area. (Note report.) This information is displayed to you in the form of a probability statement, i.e., the probability of enemy presence is equal to .80. Each problem will contain this same report - that there is an 80 percent chance that a submarine is present.

It is your task to decide, with a high degree of certainty, if an enemy submarine is, in fact, present. In order to make this decision, you will have to use your available resources to acquire information which will either <u>tend</u> to confirm or disconfirm the intelligence report. (Demonstrate). In effect, you will be "asking" the computer to report if it senses the presence of an enemy submarine.

The information which the computer displays will be unreliable, that is, any one data point, or group of data points, may be erroneous, but a large number of data points will tend to reflect the true situation.

The reliability of the information is displayed for each problem as a number between 0 and 1.00. A reliability of .75 would mean that if a submarine were present, it would be reported "present" 75 percent of the time; if it were absent, it would be reported "absent" 75 percent of the time. If reliability were only .50, then you could never correctly decide the presence or absence of a submarine more than 80 percent of the time no matter how much additional information you asked for. A reliability of 1.00 would allow you to choose with absolute certainty after acquiring only one piece of information. You can see that you would need to choose more information to compensate for a low reliability.

You may continue to ask for more data until you feel fairly certain that you can make a correct decision. When this point is reached, use the lightpen to indicate your choice. (Demonstrate.) After an intertrial interval of five seconds, the next problem will appear.

The computer will monitor your performance on each trial. If you make a decision without having evaluated enough information, a message to that effect will be displayed and the problem will be scored as an error. If you do not make a decision when you should, you will be informed which decision you should have made, and the problem will be scored as an error.

Only when you have sufficient data and make the correct decision, will you be informed that you were correct and the problem will be scored as a correct one.

The initial problems will have an associated information reliability of 1.00, meaning that you can reach a decision after only one information request. These problems will be very easy. The difficulty of the remaining problems will be selected to match your ability. This will be done because we do not want you to waste your time and effort on problems which are below your capabilities or to struggle with problems which may be too difficult.

As your performance changes with practice, the information reliability will be adjusted to insure that the problems remain challenging without being excessively difficult. You will be informed by a displayed message each time the problem difficulty is to be changed.

The procedure used to determine the sequence of problem difficulties is as follows: A running score is kept of your performance over the last five problems of the current difficulty level. If you have gotten four of them correct, then the problem difficulty will be raised. Problem difficulty will go down if you got three incorrect.

I will be in the next room monitoring your performance and will intercede if you appear to be having any difficulties. It is important that you understand what you are to do, so please ask about any sections of these instructions which may be unclear.